

WHAT IS CLAIMED IS:

- 1 1. A method for maintaining a consistent set of replicas of a database within a
2 computer cluster, comprising the steps of:
3 each node in the computer cluster receiving a database update request;
4 each node in the computer cluster voting based on a functional outcome of the
5 database update request; and
6 detecting an out-of-sync condition as a result of a different functional outcome.
- 1 2. The method as recited in claim 1, wherein the out-of-sync condition is an error.
- 1 3. The method as recited in claim 1, further comprising the step of:
2 refreshing the database in response to the detecting step.
- 1 4. The method as recited in claim 1, further comprising the step of:
2 resetting cluster membership in response to the detecting step.
- 1 5. The method as recited in claim 1, further comprising the step of:
2 blocking further participation by the node having the out-of-sync condition in
3 response to the detecting step.
- 1 6. The method as recited in claim 1, further comprising the step of:

2 declaring an end-of-transaction state on update voting completion when the
3 database update is being done in a transactional manner.

1 7. The method as recited in claim 6, further comprising the step of:
2 backing out an update when update voting does not meet a criteria established for
3 success.

1 8. The method as recited in claim 7, wherein the criteria established for success is
2 that no more than one node has inconsistent results.

1 9. A method for maintaining a consistent set of replicas of a database within a
2 computer cluster, comprising the steps of:

3 broadcasting an update to a database shared among a plurality of nodes in the
4 computer cluster;

5 applying the update to a local copy of the database at each of the plurality of
6 nodes in the computer cluster;

7 node requesting update broadcasts results of update to all of the other nodes in the
8 computer cluster;

9 comparing, by all of the other nodes in the computer cluster, the update results to
10 results of application of the update to the local copy of the database; and

11 voting, by all of the other nodes in the computer cluster, to approve update if a
12 match results from the comparison.

1 10. The method as recited in claim 9, further comprising the step of:

2 voting, by any one of the other nodes in the computer cluster, to continue with
3 update process if a match does not result from the comparison.

1 11. The method as recited in claim 9, further comprising the step of:

2 broadcasting an approval of the update to the database if all of the other nodes
3 vote to approve the update.

1 12. The method as recited in claim 10, further comprising the step of:

2 if more than one of the plurality of nodes votes to continue, performing a recovery
3 process.

1 13. The method as recited in claim 12, wherein the recovery process further comprises
2 the step of:

3 if more than a specified number of the nodes voted to continue, backing out the
4 update to the database.

1 14. The method as recited in claim 12, wherein the recovery process further comprises
2 the step of:

3 if less than a specified number of the nodes voted to continue, performing the
4 recovery process on the specified number of the nodes.

1 15. A computer cluster operable for maintaining a consistent set of replicas of a
2 database within the computer cluster, comprising:
3 a group services client operable for broadcasting an update to a database shared
4 among a plurality of nodes in the computer cluster;
5 the plurality of nodes coupled to the computer cluster operable for applying the
6 update to a local copy of the database at each of the plurality of nodes in the computer
7 cluster;
8 circuitry for broadcasting results of the update to all of the other nodes in the
9 computer cluster;
10 circuitry for comparing, by all of the other nodes in the computer cluster, the
11 update results to results of application of the update to the local copy of the database; and
12 circuitry for voting, by all of the other nodes in the computer cluster, to approve |
13 update if a match results from the comparison.

1 16. The computer cluster as recited in claim 15, further comprising:
2 circuitry for voting, by any one of the other nodes in the computer cluster, to
3 continue with update process if a match does not result from the comparison.

1 17. The computer cluster as recited in claim 15, further comprising:
2 circuitry for broadcasting an approval of the update to the database if all of the
3 other nodes vote to approve the update.

1 18. The computer cluster as recited in claim 16, further comprising:
2 if more than one of the plurality of nodes votes to continue, circuitry for
3 performing a recovery process.

1 19. The computer cluster as recited in claim 18, wherein the recovery process further
2 comprises:
3 if more than a specified number of the nodes voted to continue, circuitry for
4 backing out the update to the database.

1 20. The computer cluster as recited in claim 18, wherein the recovery process further
2 comprises:
3 if less than a specified number of the nodes voted to continue, circuitry for
4 performing the recovery process on the specified number of the nodes.

1 21. A computer program product adaptable for storage on a computer readable
2 medium, the computer program product operable for maintaining a consistent set of
3 replicas of a database within a computer cluster, comprising the program steps of:
4 broadcasting an update to a database shared among a plurality of nodes in the
5 computer cluster;
6 applying the update to a local copy of the database at each of the plurality of
7 nodes in the computer cluster;
8 node requesting update broadcasts results of update to all of the other nodes in the
9 computer cluster;
10 comparing, by all of the other nodes in the computer cluster, the update results to
11 results of application of the update to the local copy of the database;
12 voting, by all of the other nodes in the computer cluster, to approve update if a
13 match results from the comparison; and
14 voting, by any one of the other nodes in the computer cluster, to continue with
15 update process if a match does not result from the comparison.

1 22. The computer program product as recited in claim 21, further comprising the
2 program step of:
3 broadcasting an approval of the update to the database if all of the other nodes
4 vote to approve the update.

1 23. The computer program product as recited in claim 22, further comprising the
2 program step of:

3 if more than one of the plurality of nodes votes to continue, performing a recovery
4 process.

1 24. The computer program product as recited in claim 23, wherein the recovery
2 process further comprises the program step of:

3 if more than a specified number of the nodes voted to continue, backing out the
4 update to the database.

1 25. The computer program product as recited in claim 24, wherein the recovery
2 process further comprises the program step of:

3 if less than a specified number of the nodes voted to continue, performing the
4 recovery process on the specified number of the nodes.

5 A method for maintaining a consistent set of replicas of a database within a computer
6 cluster, comprising the steps of:

7 broadcasting an update to a database shared among a plurality of nodes in the
8 computer cluster;

9 applying the update to a local copy of the database at each of the plurality of
10 nodes in the computer cluster;

11 node requesting update broadcasts results of update to all of the other nodes in the
12 computer cluster;

13 comparing, by all of the other nodes in the computer cluster, the update results to
14 results of application of the update to the local copy of the database; and
15 voting, by all of the other nodes in the computer cluster, to approve update if a
16 match results from the comparison.

1 26. The method as recited in claim 9, further comprising the step of:
2 voting, by any one of the other nodes in the computer cluster, to continue with
3 update process if a match does not result from the comparison.

1 27. The method as recited in claim 9, further comprising the step of:
2 broadcasting an approval of the update to the database if all of the other nodes
3 vote to approve the update.

1 28. The method as recited in claim 10, further comprising the step of:
2 if more than one of the plurality of nodes votes to continue, performing a recovery
3 process.

1 29. The method as recited in claim 12, wherein the recovery process further comprises
2 the step of:
3 if more than a specified number of the nodes voted to continue, backing out the
4 update to the database.

- 1 30. The method as recited in claim 12, wherein the recovery process further comprises
2 the step of: ~~A~~
3 if less than a specified number of the nodes voted to continue, performing the
4 recovery process on the specified number of the nodes.

667667-20666666